Response-Time-Optimized Distributed Cloud Resource Allocation

Matthias Keller and Holger Karl

Abstract—A current trend in networking and cloud computing is to provide compute resources widely distributed exemplified by initiatives like Network Function Virtualization. This paves the way for a widespread service deployment and can improve service quality; a nearby server can reduce the user-perceived response times. But always using the nearest server is a bad decision if that server is already highly utilized.

This paper investigates the optimal assignment of users to distributed resources – a convex capacitated facility location problem with integrated queuing systems. We determine the response times depending on the number of used resources. This enables service providers to balance between resource costs and the corresponding service quality. We also present a linear problem reformulation showing small optimality gaps and faster solving times; this speed-up enables a swift reaction to demand changes. Finally, we compare solutions by either considering or ignoring queuing systems and discuss the response time reduction by using the more complex model. Our investigations are backed by large-scale numerical evaluations.

Index Terms—cloud computing; virtual network function; network function virtualization; resource management; placement; facility location; queueing model; linearisation; optimization

arXiv:1601.06262v2 [cs.NI] 29 May 2016

1 Introduction

1.1 Challenges in Distributed Clouds

A current trend in networking and cloud computing is to provide compute resources widely distributed. Computation will not only take place on desktops or large data centres, but also at smaller centres or within the network itself, e.g., inside individual in-network server racks located near backbone routers. This trend is known under different labels, for example, Carrier Clouds [6], [13], [51], Distributed Cloud Computing [3], [15], [21], [46], or In-Network Clouds [28], [48], [50]. These In-Network Clouds tend to be less cost-efficient than conventional Clouds due to a worse economy of scale; they are hence often geared towards specific network services (e.g. firewalls, load balancers). Easing a more flexible deployment of these services became popular as Network Function Virtualization [22] not only inside a data centre but also beyond, in wide area networks [17], [51], [52]. We consider not only executing network functions but more generically executing applications at those In-Network Clouds yielding an important advantage [4], [43]: The resources of these Clouds are closer to end users than those of conventional Clouds, have smaller latency between user and cloud resource, and are therefore suitable for running highly interactive applications. Examples for such applications are latency-critical applications [7], [54], user-customized streaming services [8], [18], [30], or Cloud Gaming [37]; the computing tasks range from processing the request, aggregating incoming data streams, up to rendering and encoding video streams. In such applications, the crucial quality metric is the user-perceived *response time* to a request as the application need to quickly react on user interactions. Large response times impede usability, increase user frustration [16], [37], or prevent commercial success.

An obvious solution to provide small response times would be to deploy an application at many sites so that each user finds one site nearby. This, however, is infeasible as each utilized site incurs additional costs. We are hence faced with the task to decide where a user's request shall be processed, using as few sites as possible at a best possible response time. We refer to this task as the *assignment problem*. This problem's trade-off between cost and quality is intuitive yet difficult to capture in a concrete problem statement and solution.

This difficulty lies in the nature of the response time. It is a sum of three parts:

- The network latency taken to send the request from the user to the cloud resource and sending the answer back – the *round trip time* (RTT);
- the actual *processing time* (PT) of the request;
- the *queuing delay* (QD) a request incurs at the cloud resource while other requests are currently processed at that resource (**??**).

In many applications, we can consider the processing time to be significantly smaller than the round trip time. The round trip time depends on the choice where a user's request is processed and, to a much smaller degree, on the network load along the way. The queuing delay, however, depends on the sharing of a resource among many users and is not an effect immediately influenced by the decision for a single user; it depends on the joint decision for all users.

From queuing theory we know that, for a fixed utilization, the queuing delay is shorter for higher service rates. **??** shows queuing delays for different levels of system utilization $\rho = \lambda/\mu$ for three different service rates μ. For instance, web servers answering simple requests have high service rates

- *Manuscript received January xx,yy; revised January xx, yy. This work was partially supported by the German Research Foundation (DFG) within the Collaborative Research Centre "On-The-Fly Computing" (SFB 901).*
- *The authors are with the University of Paderborn, Warburger Str. 100, 33098 Paderborn, Germany. E-mail: mkeller@upb.de, holger.karl@upb.de*

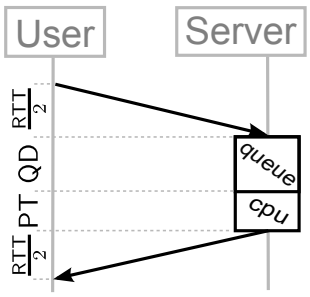

Figure 1: Response time as the sum of round trip time RTT, queuing delay QD, processing time PT.

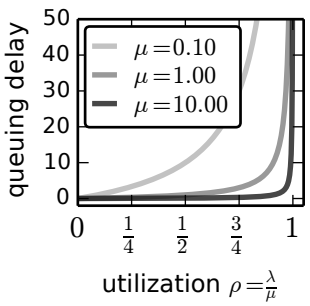

Figure 2: Queuing delays QDs for different service rates.

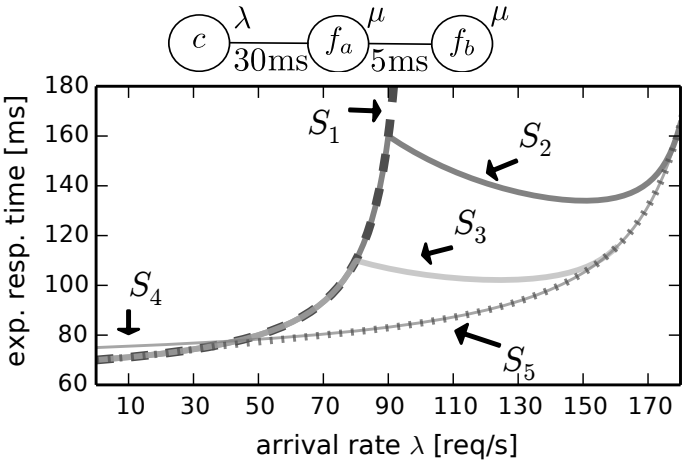

Figure 3: The lower plot shows response times of different arrival rates λ_c for three different strategies for the topology. The upper plot shows the used topology.

and often negligible queuing delays (say, below $10\,\mathrm{ms}$). This may explains why they are commonly ignored in literature (**??**). We focus in this paper on computation-intensive applications, such as data processing, intrusion detection, or game render application (see next paragraph) oppose to light-weight application with nearly no computation like web servers just returning static content. The processing time of computation-intensive applications is longer (e.g. up to $1\,\mathrm{s}$) than of light-weight applications, which implies a lower service rate and a longer queuing delay. A long queuing delay becomes a large portion of the response time, significant large enough to necessitates considering them when deciding the assignments.

Shah et al. [49] survey intrusion detection systems and cite different measurements of packet processing times of up to $10\,\mathrm{ms}$. Barker et al. [7] study game server map loads lasting 20–$110\,\mathrm{ms}$ in their experiments. Ishii et al. [30] conduct experiments on AWS [1] using a parallel Data Processing Application and observe processing time between 400 and $1800\,\mathrm{ms}$. Lee et al. [37] and Claypool et al. [16] observe drop in user experience when playing computer games with artificially increased latency to larger than $200\,\mathrm{ms}$. In summary, we focus on applications with long processing times, $10\,\mathrm{ms}-1\,\mathrm{s}$, and on average internet round trip times, 60–$600\,\mathrm{ms}$ [36].

Different proportions between round trip times and processing times are possible. Deciding the assignment in such scenarios can be simplified by ignoring the less dominant part: Very short round trip times inside data centres, say less than $20\,\mathrm{ms}$, leaves queuing delays as the dominant part of the response time. Similarly, long processing times, say more than $1\,\mathrm{min}$, renders queuing delays as the dominant part. In both cases, round trip times can be ignored; doing so, the assignment problem becomes a simpler mapping problem. On the opposite site, very short processing times, say $0.1\,\mathrm{ms}$[1], result in very low queuing delays rendering the round trip time being the dominant part. When ignoring queuing delays in this case, the assignment problem becomes a simpler, non-convex Facility Location Problem. In summary, only if round trip times and processing times are not of the same magnitude, dropping the less dominant part is a vital option. If both times are of the same magnitude, also queuing delays have to be considered when deciding the assignment. Ignoring queuing delays in such cases worsen the assignment resulting in high response times; we compare both cases with different proportions of round trip times and processing times (**??**).

1.2 Queuing Delay Effects

How does the queuing delay affect the response time? As a toy example, let us consider the network from **??** with three locations of interest: One client c and two possible facility locations f_a, f_b with compute resources to run the application. These resources are equally fast and can serve requests at rate $\mu{=}100\,\mathrm{req./s}$. Assume the round trip times between c and f_a as $60\,\mathrm{ms}$ and between c and f_b as $70\,\mathrm{ms}$. Requests enter the network at c with arrival rate λ. With this setup, the requests can be served at only f_a, only f_b, or split among f_a and f_b. **??** shows, as a function of the arrival rate, the resulting response times RT = RTT + PT + QD for a few simple strategies S_i[2].

The first strategy S_1 minimizes only the requests' round trip time: Requests are assigned to the nearest facility f_a and if its capacity μ is exceeded the remaining requests are assigned to f_b. The dramatic response time growth for $\lambda{>}80$ is the result of too many requests assigned to facility f_a. Let λ_a be the assigned requests to f_a, then f_a's utilization ρ_a is λ_a/μ. To avoid too large utilizations, strategies S_2 and S_3 limit them to a maximum value $\hat{\rho}$, $\rho_a,\ \rho_b \leq \hat{\rho}{<}1$; S_2 uses $\hat{\rho}{=}0.9$ and S_3 uses $\hat{\rho}{=}0.8$. On the one hand S_3 with a lower limit has a shorter RT (**??**) than S_2 as the second facility is used earlier. But on the other hand S_2 can handle a higher arrival rates than S_3, $160{<}\lambda{<}180$, because a higher limit enables handling more requests in total; system capacity is $\lambda{<}2\mu\hat{\rho}$. To relax such a predefined upper bound, S_4 dynamically adjusts the limit to the current system utilization, $\hat{\rho}{=}\lambda/2\mu$. With the same resources at both locations S_4 boils down to evenly splitting the load between the two facilities. Compared to $S_{1..3}$ the resulting RTs are on the one hand small for $\lambda{>}40$ but on the other hand larger for $\lambda{<}40$.

So far, all assignment strategies ignore the resulting queuing delays. In contrast, our last strategy S_5 additionally

1. Corresponding service rate of $10{,}000\,\mathrm{req./s}$ of a high performance webserver, e.g. Apache, Nginx, delivering static content.

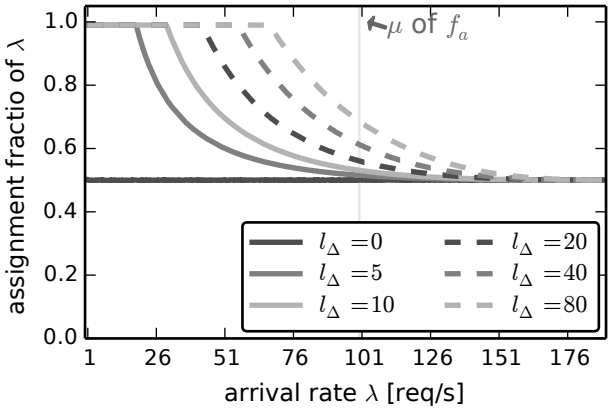

Figure 4: Request assignments to f_a with different distances l_Δ between f_a and f_b.

reduces the queuing delays of both resources. The strategy's request assignment depends on the round trip times (l_{cf}) between both resources. The resulting RTs are the lowest for all strategies $S_{1..5}$. In conclusion, we were able to improve assignments by considering queuing delays.

Hereafter, we list the equations of the expected response times in **??**: S_1 to S_3 results in (1), S_4 in (2), and S_5 in (3).

$$f_{\mu,\hat{\rho}}(\lambda) := \begin{cases} 60 + \frac{1}{\mu-\lambda}, & \text{if } \lambda \leq \hat{\rho}\mu \\ \frac{\hat{\rho}\mu}{\lambda}\left(60 + \frac{1}{\mu-\hat{\rho}\mu}\right) + & \\ \quad \frac{\lambda-\hat{\rho}\mu}{\lambda}\left(70 + \frac{1}{\mu-\lambda+\hat{\rho}\mu}\right), & \text{else} \end{cases} \quad (1)$$

$$g_\mu(\lambda) := \frac{1}{2}\left(60+\frac{1}{\mu-\lambda/2}\right) + \frac{1}{2}\left(70+\frac{1}{\mu-\lambda/2}\right) \quad (2)$$

$$h_\mu(\lambda) := \min_{\lambda_1 \in [0,\,\lambda]} \left\{ \begin{array}{l} \frac{\lambda_1}{\lambda}\left(60 + \frac{1}{\mu-\lambda_1}\right) + \\ \frac{\lambda-\lambda_1}{\lambda}\left(70 + \frac{1}{\mu-\lambda+\lambda_1}\right) \end{array} \right\} \quad (3)$$

?? shows strategy S_5's request assignments to resource f_a as a fraction of λ on the vertical axis; the remaining requests are assigned to f_b. The horizontal axis shows an increasing arrival rate λ. The different lines correspond to distances l_Δ between f_a and f_b – how much longer request transportation takes to send to f_b instead to f_a. This way the original toy example is line l_Δ=10 and the other lines vary the round trip time to f_b. If the resources have the same round trip time, the assignment results in an even split. However, if one resource is farer away, l_Δ>0, at first the nearer resource is preferred and only with increasing arrivals do the assignments converge to an even split. Then, the queuing delay portion of the response time is significantly larger than l_Δ.

1.3 Contribution

This paper discusses finding the optimal assignment between requesting users and compute resources hosting the answering application at different locations in the network. The assignment minimizes the expected average response time for all users. We claim the necessity of considering request's queuing delays at used compute resources to avoid suboptimal assignments to e.g. over-utilized resources (when round trip times and queuing delays are of the same magnitude). To proof this claim, we present an extended Facility Location Problem with integrated queuing systems (**??**), show its convexity, and, for the first time, obtain optimal[3] solutions for larger networks using a convex solver (**??**). Due to problem's complexity, solving times were already large for medium-sized networks (e.g. **??** `ta2` topology) hindering the large scaled evaluation we had envisioned. We were able to shorten the solving times with high accuracy by non-trivially linearising the convex problem (**??**). This linearised problem has a larger search space as the convex problem; despite this, it solves the original problem significant faster (empirically shown, **??**). Having now an adequate and fast substitution at hand enables us to compare numerous solutions obtained by considering and ignoring queuing delays supporting our claim by showing significant response time reduction when considering queuing delays (**??**). In addition, we show how the response time and queuing delays increases when using lesser compute resources (e.g. fewer locations, **??**).

This evaluation in **??** extends our own previous work [33]. We compare four factors influencing queuing delays and in addition vary input randomly in order to verify the statistical relevance of our findings. In summary, we solved and analysed $52{,}500$ configurations.

2 Related Work

Assignment problems of the form described above have been investigated before. We structure their comparison along four dimensions relevant to this work: Their model complexity, simplifications reducing the problem's search space, optimization goals, and solution approaches. Finally, related systems of geographical load balancing are compared.

2.1 Model Complexity

The simplest model considers *only the round trip time* (RTT) when assigning users to cloud resources. They equate response time with RTT. Clearly, this is a simplification of reality, yet minimizing this average RTT is equivalent to the well-known capacitated Facility Location Problem (FLP). If the problem is further restricted to only use p resources, it becomes a p-median FLP, which is NP-hard [31].

A step closer to reality is modelling *also the processing time* (PT) in addition to the RTT. But as long as PT is constant, this still stays a Facility Location Problem of the type described above. This can be easily seen by extending the original network topology by pseudo-links (at the server or user side) that represent these processing times via their latencies; this is a common rewriting technique for graph-based problems (including Facility Location Problems).

The real challenge occurs when we also *consider the queuing delay* (QD). In this case, the additional time cannot be expressed by rewriting the network topology as the QD depends on the assignment decisions: A higher utilization results in a longer wait, possibly trading off against a shorter RTT.

So far, this more general model has been considered only by few works discussed in the remaining of this section, most use simpler assumptions than ours (**??**) rendering the

3. Numerically obtaining solutions by solver with a gap threshold of 10^{-6}.

problem easier to solve. Vidyarthi et al. [53] allow the same degrees of freedom as we do. They approximate, similar to us, the non-linear part of the objective function with a piece-wise linear function. However, in contrast to our work, they used a cutting plane technique which iteratively refines the piece-wise function as necessary; it remains unclear how large their linearisation error is. In contrast, our evaluation shows small linearisation errors; and this is achieved by using a simpler technique.

2.2 Simplifications

Other authors investigate slightly different scenarios, so that their problem formulations are similar, yet simpler than ours.

Some authors [55], [59] replace the non-linear QD part with a constant upper bound and, consequently, the resulting problems become simpler to solve. But this also hides QD changes as a result of assignment changes. For instance, in a situation where load balancing would reduce the QD, this reduction is not visible as the QD part is constant. Consequently, the resulting solution has further potential for optimization – we exploit this potential.

In another simplification, the assignments are predefined by a rule. Some authors [2], [55], [59] always assign requests to the nearest cloud resource. In such a case, the problem reduces to just finding the best resource location and is easier to solve. The assignments are then predetermined by the rule. However, balancing the assignments could further reduce the QD but is not considered. We do not use any predefined assignment rule, so we have the freedom to change assignments in order to further reduce the response times.

Another group of authors [10], [19] uses a parametrized assignment rule called the gravity rule: Weights determine how users are assigned to cloud resources. These configurable weights are used to continuously solve the same problem with new weights reflecting the resource utilizations of the previous solution. This approach does not guarantee to converge, so the authors propose a heuristic that attenuates the changes in each iteration, enforcing convergence with an unknown linearisation error. In contrast, we solve the problem in one step by using all information to find the global optimum.

Liu et al. [39], Lin et al. [38], and Goudarzi et al. [25] present a similar Facility Location Problem with convex costs such as queuing delays or resource's energy costs. In contrast to our work, they relax the integer allocation decision variable simplifying the problem to the cost of a less accurate solution when rounding up the obtained continuous allocations. Our goal, in contrast, is to prevent unexpected expenses by introducing an upper bound to the number of used resources. Continuously relaxing our problem can cause any location to be allocated a bit and, consequently, any site is used and paid. While the papers [38], [39] only consider queuing delays as a cost function, this paper discusses a holistic queuing system integration and additionally considers splitting and joining (assigning) of the arrival process.

2.3 Optimization Goal

Existing literature uses queuing delays in FLPs with three optimization goals: classic FLP, min/max FLP, and coverage FLP.

Classic FLPs are problems that minimize the average response time, like our problem (**??**) or others [10], [19], [47], [53], [55], [59], allowing RT variations for individual users.

Aboolian et al. [2]'s min/max problem minimizes the maximum response time. Intuitively, such problems improve especially the users' RT with high RTTs to cloud resources. However, if only one such user exists with resources being far away, assigning this user will negatively affect the assignments of other users: Their assignments are now less-restrictedly constrained by a relaxed upper bound and are likely worse than without the first user. In contrast, classical FLPs do not suffer this way from a worse case user.

Another type of problem is coverage problems; the user assignment's response times is upper bounded [40], [41], [42]. Structurally, a coverage problem is a special, simpler case of a min/max problem; the first has a predefined bound, which is additionally minimized in the second. Intuitively, such problems can be applied in scenarios where service guarantees for a certain maximal response time will be provided and paid. In contrast, classical FLPs allow minimizing the average response time below the lowest possible response time bound.

2.4 Solution Approaches

A couple of heuristics were proposed solving related problems which are variants of the NP-hard capacitated FLP [27]. No work so far used solvers to obtain solutions (for non-relaxed problems) and full enumerations are known for small instances limited to open five facilities [2]. A greedy *dropping* heuristic successively removes from the set of candidates that resource which increases the response time by the smallest amount [55]. Greedy *adding* heuristics successively add resources, which decreases the response time by the largest amount [2], [10], [19]. Another heuristic probabilistically selects set changes of used resources [19] or performs a breath-first-search through "neighbouring solutions" where two solutions are neighbours if their sets of used resources differ in one element [2], [19]. Such heuristics can be stocked in local optima and to mitigate this drawback meta-heuristics are used as a superstructure [2], [10], [19], [47], [55]. These meta-heuristics typically refine previously generated initial solutions, which are obtained randomly or by combining existing solutions. The hope is that among the found local optima, one solution is very close to the global optimum – but without any guarantee. In contrast, we obtain global optima. This is an important step for heuristic development as only this enables a clear judgement of heuristics' accuracy; their solution's gap to the global optimum.

Others [53], [55] may achieve near optimal solutions by using optimization techniques like branch-and-bound and cutting planes but their solutions have unknown optimality gabs. In summary, either optima for small input or solutions with unknown optimality gap are obtained. This motivated our work on finding near-optimal solutions with a numerically very small optimality gap.

Liu et al. [39] and Wendell et al. [56] present distributed algorithms for their global Geographical Load Balancing problem by decomposing it into separate subproblems solved by all clients. These subproblems converge to the optimal solution only if they are executed in several synchronized rounds in which assignment and utilization information are exchanged among all clients. Both papers state that this distributed algorithms would obtain optimal solution faster than gathering everything to a centralised solver. However, we believe that each round a communication delay is introduced when sending update information among all clients; they had ignored this delay in their evaluations. The resulting total delay over all rounds is likely larger than communicating with a centralised coordinator. In addition, our p-median Facility Location Problem has a global constraint on the maximal used resources preventing it to be easily separated into subproblems.

We observed that problem instances were solved only exemplary so far [2], [10], [19], [38], [39], [41], [47], [47], [56]. Consequently, the average performance of these solution approaches is hard to predict. We go beyond this by undertaking a statistical performance evaluation. We randomly vary our input data and verify the statistical relevance of our findings.

2.5 Geographical Load Balancing

A system for Geographical Load Balancing (GLB) comprises two parts: The decision part selects appropriate server, sites, or Virtual Machines for requests of a certain origin – the previous sections considers them. This section focuses on the realisation part, which gathers monitoring information and implements selections. Different middlewares had been proposed [23], [56], [57], [58] which are shared between applications. In this way, each application benefits from instances of the other application running at diverse sites by sharing monitoring information such as latency to servers or to customers. They realise request assignments, e.g., to close-by or low utilised server by either configuring the Domain Name System (DNS) or are explicitly queried ahead a request send. Slightly different, Cardellini et al. [14] propose redirecting requests to different sites to balance the load. Policies ranges from redirecting all, only largest, or only group requests to selecting sites based on round-robin, site utilization, or connection properties. Our paper focuses on solving the problem and investigates whether the complexer problem with queuing systems is worth the additional efforts and our results can be applied to improve geographical load balancing systems.

3 Problem

This section first formalises our scenario model and then details on practical realisations. Afterwards this section discusses problem's convexity and proposes a problem linearisation minimizing the maximal linearisation error.

3.1 Model

Our scenario is formalized as a capacitated p-median Facility Location Problem [20]. A bipartite graph $G = (C \cup F,\ E)$ has two types of nodes: *clients* ($c \in C$) and *facilities* ($f \in F$). Clients correspond to locations where user request flows enter the network. Facilities represent candidate locations to execute the application, e.g., data centres. More precisely, a (compute) resource refer to a host at such a data centre executing the application. **??**a shows such a graph. The geographically[4] distributed demand is modelled by the request arrival rate λ_c for each client c. Computing capacity is modelled as the request serving rate μ_f for each facility f. The round trip time l_{cf} is the time to send data from c to f and back. **??** lists all variables.

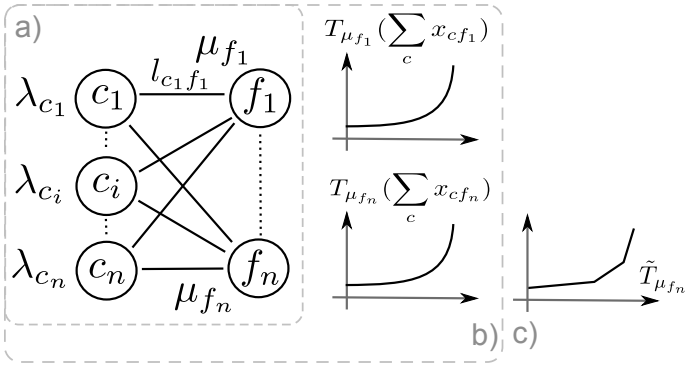

Figure 5: Bipartite graph of a Facility Location Problem (a); time-in-system functions at each facility (b) and, alternatively, piece-wise linearised functions (c).

Table 1: Model variables

Input constants:	
$G = (V, E)$	Bipartite graph with $V = C \cup F$, $C \cap F = \varnothing$ with client nodes $c \in C$ and facility nodes $f \in F$
$l_{cf} \in \mathbb{R}_{>0}$	Round trip time between c and f
$\mu_f \in \mathbb{R}_{>0}$	Service rate as capacity at f
$\lambda_c \in \mathbb{R}_{>0}$	Arrival rate as demand at c
$\mathsf{T}_\mu \in \mathbb{R}_{>0}$	Time in queuing system (TiS)
$\alpha_{\mu s}, \beta_{\mu s}$	s-th basepoint $\mathsf{T}_\mu(\alpha_s)=\beta_s$ of $\tilde{\mathsf{T}}_\mu$
Decision variables:	
$x_{cf} \in \mathbb{R}_{>0}$	Assignment in demand units
$y_f \in \{0, 1\}$	Indicator if f is opened ($= 1$)
$z_{fs} \in [0, 1]$	Weight of s-th basepoint at f
Helper variables:	
Λ; Λ_f	total $\Lambda = \sum_c \lambda_c$; at f: $\Lambda_f = \sum_c x_{cf}$
$\tau \in \mathbb{R}_{>0}$	Sufficient small value
$\rho \in \mathbb{R}_{>0} = \lambda/\mu$	System utilization

Our first problem formulation recapitulates the known p-median problem $\mathsf{P}(G,\ \lambda,\ \mu,\ p)$:

$$\min_x \quad \frac{1}{\sum_c \lambda_c} \sum_c \sum_f x_{cf} l_{cf} \qquad \text{(objective)} \qquad (4)$$

$$\text{s.t.} \quad \sum_f x_{cf} = \lambda_c, \qquad \forall c \qquad \text{(demand)} \qquad (5)$$

$$\sum_c x_{cf} \le y_f \mu_f, \qquad \forall f \qquad \text{(capacity)} \qquad (6)$$

4. More precisely, the request arrival and service points are topologically distributed; the round trip time of a path between two points only roughly matches its geographically distance. We use "geographically" for a convenient explanation.

$$\sum_f y_f = p \qquad \text{(limit)} \qquad (7)$$

The formulation contains two decision variables: $x_{cf} \in \mathbb{R}_{\geq 0}$ describes which part of c's request rate λ_c is assigned to which f; $y_f \in \{0, 1\}$ describes if location f is used or not. The objective is to minimize the average response time; but without modelling the queuing delay and service time at facilities, the response time only consists of the round trip time. The RTT is minimized while all demand is served (5) and the capacity is not exceeded (6).

In addition, exactly p locations are used (7). This constraint serves two proposes. First, by limiting the number of location where the application is developed to, the expenses for the application provider when leasing Cloud resources is bound. In Facility Location variants where facility opening costs are directly integrated the resulting total costs are unsure. Second, stating the problem with this bound allow us to investigate the response time trend while allowing more and more resources (**??**). Since 1979 the problem without capacity is known to be NP-hard [27]. This problem is a generalization and, thus, also NP-hard.

Until now, the response time has only been the round trip time. To predict the queuing times, the model is extended by queuing systems at each facility (**??**b). There, the service times are exponential distributed. The inter-arrival times at each node c are described by a Poisson process. The requests can be assigned to multiple facility ($\sum_f x_{cf}$) and, there, the individual assignment from different nodes are aggregated ($\sum_c x_{cf}$). The resulting process is also a Poisson process, because splitting and joining does not change the underlying random distribution. As a result, we have a M/M/1-queuing model [11]. The function for the time in queuing system (TiS) computes the processing time plus the queuing delay (**??**), $\mathsf{T}_\mu(\lambda) = \frac{1}{\mu - \lambda}$. Putting everything together, the corresponding formulation of this queuing-extended p-median problem $\mathsf{QP}(G, \lambda, \mu, p)$ is:

$$\min_{x,\,y} \quad \underbrace{\frac{\sum_{cf} x_{cf} l_{cf}}{\sum_c \lambda_c}}_{\text{average RTT}} + \underbrace{\frac{\sum_f \left(\sum_c x_{cf}\right) \frac{1}{\mu_f - \sum_c x_{cf}}}{\sum_c \lambda_c}}_{\text{average TiS}} \qquad (8)$$

$$\text{s.t.} \quad \sum_f x_{cf} = \lambda_c, \quad \forall c \quad \text{(demand)} \qquad (9)$$

$$\sum_c x_{cf} < y_f \mu_f, \quad \forall f \quad \text{(capacity)} \qquad (10)$$

$$\sum_f y_f = p \qquad \text{(limit)} \qquad (11)$$

The new objective (8) is to minimize the average response time, which is the sum of the average round trip time and the average time in system (**??**). Constraint (9) is the same as Constraint (5); all demand must be served. Constraint (10) assures the *steady state* ($\lambda < \mu$, c.f. [11]) for each queuing system. Finally, Constraint (11) mandates to use exactly p locations, just like Constraint (7).

3.2 System design

The presented optimisation problem QP is part of a large system which dispatches requests of a certain origin to sites as decided. Examples of such a system ranges from Geographical Load Balancing systems (**??**) to our own Application Deployment Toolkit [34]. They monitor traffic, decide assignments, and reconfigure the dispatching subsystem in time periods. The average arrival rate λ_c is the averaged number of incoming requests at router c for the last period. By solving problem QP once a period, the request assignments[5] are decided for the next period. The decision is realised by configuring the dispatching subsystem, e.g., DNS, and allocating cloud resources accordingly. The system is not meant to allow a fine grained assignment decision for each incoming request, e.g. at line speed. On longer terms, it decides which sites are strategically used and how incoming requests are roughly distributed.

3.3 Convex Optimization

Previous work (**??**) also considered our objective function (8) but did not solve the corresponding problem optimally, except for small graphs via full enumeration. This is because of the non-linearity of the objective function which necessitates non-linear solvers. There exist a couple of non-linear solvers with different specializations: quadratic, convex, or non-convex objective functions. By determining the complexity class of our objective function, we can choose a suitable solver, to efficiently obtain a global optimum.

This section first proves the objective function's convexity and shows that it is not simpler, e.g., quadratic. Afterwards, it describes how we used the convex solver.

Definition 1. A function g is convex if its domain $\mathrm{dom}(g)$ is a convex set and if $g''(x) \geq 0$ holds $\forall x \in \mathrm{dom}(g)$ [12].

Lemma 1. Function $g = \sum_i w_i g_i$, $g: \mathbb{R}^n \to \mathbb{R}$, is convex, iff $\forall i$: $g_i: \mathbb{R}^n \to \mathbb{R}$ and $w_i \in \mathbb{R}_{>0}$ is convex [12].

Theorem 1. The objective function (8) of QP is convex with function T_μ computing the sojourn time in an M/M/1 queuing system.

Proof: The domain of $\mathsf{T}_\mu(\lambda) = 1/\mu - \lambda$ is the interval $0 \leq \lambda < \mu$ enforced by constraint (10); an interval is always a convex set. The second derivative $\mathsf{T}''_\mu(\lambda) = 2/(\mu - \lambda)^3$ is always larger 0 within its domain. By **??**, $\mathsf{T}_\mu(\lambda)$ is a convex function.

For a fixed f in the objective function, $0 < \Lambda_f = \lambda < \mu$ and $\mathsf{T}_\mu(\Lambda_f)$ is convex. Then, the non-negative weighted sum of convex functions $\sum_f \Lambda_f \mathsf{T}_\mu(\Lambda_f)$, $\Lambda_f = \sum_c x_{cf}$ is also convex (**??**). The term remains convex after $1/\Lambda > 0$ is multiplied. The left term of the objective function is linear and also convex. Since the sum of two convex functions is convex, the objective function (8) is convex. □

With the knowledge of a convex objective function, we can ignore less efficient solvers for more general, non-convex problems. The next more efficient solver class is quadratic, which need objective functions of the form $x^T M x$ with symmetric matrix $M \in \mathbb{R}^n$. But our objective function is not of this form, making quadratic solvers inapplicable. Consequently, we have to use a convex solver.

Implementation: We choose the optimization framework CVXOPT [5] from the authors of [12]. QP is a mixed integer problem, which is not directly supported by CVXOPT. Continuously relaxing the problem is not possible (**??**). We

5. The assignment x_{cf} is the *request rate* dispatched from c to f, in short *request assignment*.

decomposed QP into solving multiple subsets F' of F with $|F'| = p$:

$$\mathsf{QP}(G=(C\cup F,E),\,\lambda,\,\mu,\,p,\,\tau) = \min_{F'\in F,\,|F'|=p}\{\,\mathsf{PQP}((C\cup F',E),\,\lambda,\,\mu,\,\tau)\,\} \quad (12)$$

with the purely convex sub-problem $\mathsf{PQP}(G,\,\lambda_c,\,\mu_f,\,\tau)$:

$$\min_x \quad \frac{\sum_{cf} x_{cf} l_{cf}}{\sum_c \lambda_c} + \frac{\sum_f \frac{\sum_c x_{cf}}{\mu_f - \sum_c x_{cf}}}{\sum_c \lambda_c} \quad (13)$$

$$\text{s.t.} \quad \sum_f x_{cf} = \lambda_c, \qquad \forall c \quad \text{(demand)} \quad (14)$$

$$\sum_c x_{cf} \le \mu_f - \tau, \quad \forall f \quad \text{(capacity)} \quad (15)$$

The decomposition optimally solves QP by solving a non-integer convex subproblem PQP several times for different configurations of the binary variables y_q; these variables indicate which facility is used. First, problem PQP is solved with all facility subsets $Q'\subseteq Q$ with $|Q'|=t$. Then, one of all problem PQP's solution is selected that has the minimal response time (13). In this solution, the decision vector x equals the QP's decision vector x and problem QP's decision vector y is represented by subset Q', $\forall q\in Q'$: $y_q=1$. In this way, the optimal solution for problem QP is found.

CVXOPT solves PQP by checking the domain (constraints) and iterating towards the optimum by using the Jacobi and Hessian matrix (first and second order derivatives) of the objective function (13). Hardcoding such matrices is not feasible for a large number of parameter configurations. We want to have an automated solution obtaining these matrices at runtime. Algebra systems like Maxima[6] can be used, but need a detour through another system and computing derivatives of multi-dimensional functions takes time; for small input, more time than solving the problem. To obtain these matrices faster, we found, not too surprisingly, that the structure of (13) and its derivatives are the same for different $|C|$, $|F|$.

Exploiting this property, we were able to deduce a construction rule for both matrices. Using this rule, we construct our Jacobi and Hessian matrices at runtime for different inputs without notable overhead.

In detail, we constructed the Jacobi and Hessian matrices from the objective function (13); here, reintroduced as a convenient copy f (16), $f(x_{11},...,x_{cf})$.

$$f(x) = \frac{1}{\Lambda}\sum_{cf} l_{cf}x_{cf} + \frac{1}{\Lambda}\sum_f \frac{\Lambda_f}{\Lambda_f-\mu_f}, \qquad \Lambda = \sum_c \lambda_c,\ \forall f\colon \Lambda_f = \sum_c x_{cf} \quad (16)$$

The Jacobi matrix (17) for one function is a vector of partial derivatives $(a_{11},...,a_{cf})$ for each variable x_{cf}.

$$J_f(x) = (a_{cf})_{cf\in C\times F}, \qquad a_{cf} = \frac{\Lambda_f}{\Lambda(\mu_f-\Lambda_f)^2} + \frac{1}{\Lambda(\mu_f-\Lambda_f)} + \sum_c \frac{l_{cf}}{\Lambda} \quad (17)$$

Structurally, $f(x)$ is a sum of terms, and differentiating $f(x)$ can be done by differentiating the terms individually and afterwards summing all terms up. Two types of terms exist (18) with different derivatives (19, 20). In Jacobi matrix (17), each partial derivative a_{cf} is $g'_{1,cf}(x) + g'_{2,cf}(x)$.

$$g_1(x) = \frac{l_{cf}x_{cf}}{\Lambda} \quad g_2(x) = \frac{\Lambda_f}{\Lambda(\Lambda_f-\mu_f)} \quad (18)$$

$$g'_{1,cf}(x) = \frac{d}{dx_{cf}}\frac{l_{ij}x_{ij}}{\Lambda} = \begin{cases}\frac{l_{ij}}{\Lambda} & \text{if } cf{=}ij\\ 0 & \text{else}\end{cases} \quad (19)$$

$$g'_{2,cf}(x) = \frac{d}{dx_{cf}}\frac{\Lambda_j}{\Lambda(\Lambda_j-\mu_j)} = \begin{cases}\frac{\Lambda_j}{\Lambda(\mu_j-\Lambda_j)^2} + \frac{1}{\Lambda_j(\mu_j-\Lambda_j)} & \text{if } f{=}j\\ 0 & \text{else}\end{cases} \quad (20)$$

Similarly, the Hessian matrix (21) contains second-order, partial derivatives which are first derived in x_{cf} direction (rows) and then in x_{de} direction (columns).

$$H_f(x) = (a_{cf\,de})_{cf\in C\times F,\,de\in C\times F}, \qquad a_{cf\,de} = \begin{cases}\frac{2\Lambda_f}{\Lambda(\Lambda_f-\mu_f)^3} + \\ \quad \frac{2}{\Lambda(\Lambda_f-\mu_f)^2} & \text{if } f{=}e\\ 0 & \text{else}\end{cases} \quad (21)$$

Each cell $a_{cf\,de}$ (21) is $g''_{1,cf\,de}(x) + g''_{2,cf\,de}(x)$ from (22, 23).

$$g''_{1,cf\,de}(x) = \frac{d}{dx_{cf}\,dx_{de}}\frac{l_{ij}x_{ij}}{\Lambda} = 0 \quad (22)$$

$$g''_{2,cf\,de}(x) = \frac{d}{dx_{cf}\,dx_{de}}\frac{\Lambda_j}{\Lambda(\Lambda_j-\mu_j)} = \begin{cases}\frac{2\Lambda_j}{\Lambda(\Lambda_j-\mu_j)^3} + \frac{2}{\Lambda(\Lambda_j-\mu_j)^2} & \text{if } f{=}j \wedge e{=}j\\ 0 & \text{else}\end{cases} \quad (23)$$

3.4 Linear Approximation

While CVXOPT solves the problem optimally, it has to test all subsets F', which takes time. As an alternative, the convex objective function is linearised. This way, well researched linear solvers can be used to obtain solutions faster.

3.4.1 Piece-wise linear

Any non-linear function $g(x)\colon\mathbb{R}\to\mathbb{R}$ over a finite interval $[\alpha_0,\,\alpha_{m-1}]\subset\mathbb{R}$ can be approximated by a piece-wise linear (PWL) function $\tilde{g}$ [24]. This function consists of m basepoints $\alpha_0,..,\alpha_s,..,\alpha_{m-1}$, corresponding function values $\beta_s{=}g(\alpha_s)$, and is defined in (24) for $\alpha_s{\le}x{\le}\alpha_{s+1}$.

$$\tilde{g}(x) := (x-\alpha_s)\frac{(\beta_{s+1}-\beta_s)}{(\alpha_{s+1}-\alpha_s)} + \beta_s, \qquad \alpha_s \le x \le \alpha_{s+1} \quad \forall s{\in}[0,m{-}2] \quad (24)$$

As an example, let us consider the part $\lambda\mathsf{T}_\mu(\lambda)$ of (8) for $\mu{=}1.0$. Then $g(\rho){:=}\lambda\mathsf{T}_1(\lambda){=}\lambda/1-\lambda$ is our example function to linearise. **??**a shows g and two different linearisations $\tilde{g}_1$ and $\tilde{g}_2$. The horizontal axis shows the arrival rate and the vertical axis shows the corresponding TiS. **??**b shows the absolute differences between g and either linearisation $\tilde{g}_1$ or $\tilde{g}_2$. These differences denote the linearisation accuracy: The smaller the differences are, the tighter the PWL function resembles the original function. We use the maximum of all absolute differences $\epsilon_{\tilde{g}}$, defined in (25), to measure the

6. Maxima manuel: http://maxima.sourceforge.net/docs/manual/maxima.pdf

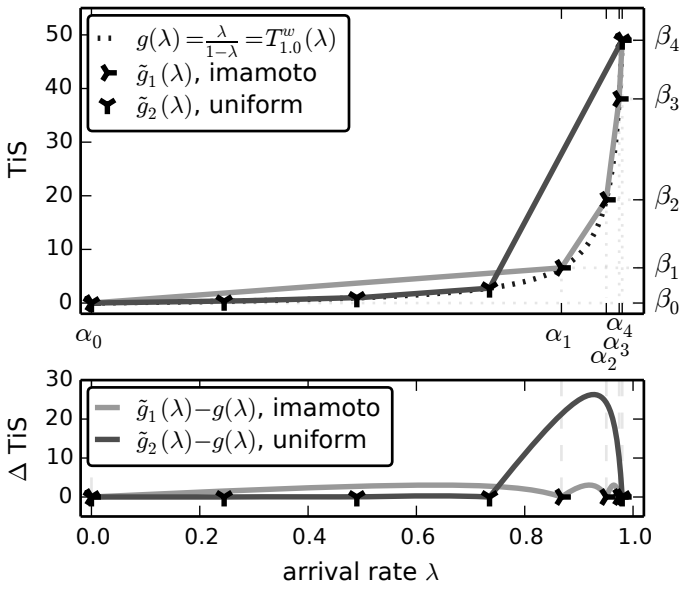

Figure 6: The top plot shows an example function g with two possible linearizations with the same number of basepoints. The bottom plot shows absolute differences between the linearizations and g. Imamoto's linearisation has a smaller maximum difference.

linearisation accuracy. We seek basepoints α_i that minimize this error.

$$\epsilon_{\tilde{g}} := \max_{x \in [\alpha_0,\, \alpha_{m-1}]} |\tilde{g}(x) - g(x)| \tag{25}$$

Given a set of basepoints resulting in a certain error, this error is reduced by placing an additional basepoint at a point where the absolute difference equals the error. However, more basepoints also increase the number of necessary variables for the optimization, which increases search space and solving runtime.

Some functions are hard to approximate with linear segments, e.g., functions with large second-order derivative values. If their values are large within the linearisation interval $[\alpha_0,\, \alpha_{m-1}]$, the error will be large. The TiS function's asymptote $\lim_{\lambda\to\mu} \mathsf{T}_\mu(\lambda)=\infty$ approximated by linear segments results in such a larger error. One possible control knob is to adjust the interval $\alpha_{m-1}<\mu$. But this also introduces an artificial capacity limit: Small values (e.g. $\alpha_{m-1}=0.8\mu$) result in fewer requests served than possible (c.f. **??**b's S_2). Consequently, the total arrival rate for which solutions are feasible to obtain is smaller, $\sum_c \lambda_c/p \leq \alpha_{m-1} < \mu$ with p used resource.

Both PWL functions in **??**b, uniform and imamoto, have the same number of basepoints but at different positions. As shown, uniformly distributing the basepoints can dramatically increase the error ($\tilde{g}_1$). In contrast, the grey basepoints have small errors ($\tilde{g}_2$). Those basepoints were computed by our algorithms detailed in **??**.

We evaluate the first two control knobs, the number of basepoints and the linearisation interval's upper bound, in **??**. For the third control knob, the basepoint positions, our algorithm determines basepoints with low error.

3.4.2 Linearisation algorithm

Our algorithm obtains basepoints for convex functions with low error. It is an extended version of Imamoto's algorithm [29]. Imamoto's algorithm iteratively refines m basepoints by moving them individually along the abscissa to reduce the error $\epsilon_{\tilde{g}}$. Each basepoint's adjustment Δ_s along the abscissa, $\alpha_s^{\text{neu}}=\alpha_s^{\text{alt}}+\Delta_s$, is computed from the basepoint's first-order derivative $\frac{d}{d\alpha_s^{\text{alt}}} g(\alpha_s^{\text{alt}})$ and the inter-basepoint distance $d_s=\alpha_s-\alpha_{s+1}$.

The paper's [29] statement is that the algorithm computes basepoints which have the maximal linearisation accuracy for the given number of used basepoints. However, the algorithm runs in numerical issues rendering the algorithm useless for some convex functions. When fixing[7] them, it cannot be guaranteed any more that the resulting basepoints form an linearisation with maximal accuracy (minimal error). But it is still very small – still a good and fast option to linearise convex functions.

More in detail, we extend Imamoto's algorithm [29] and fixed the following two cases: First, the algorithm iteratively adjusts the current set of basepoints so that the error is successively reduced. These adjustments are weighted in order to allow gradually finer changes so that the error after each iteration converges to the minimum error in theory. In practice, floating-point accuracy is limited and sometimes values are too small, changes not applied, and the algorithm iterates infinitely. We fixed that by additionally aborting if no further basepoint changes are observed.

Second, for special functions the algorithm terminates with a division by zero. The cause is computing a basepoint α_s's adjustment Δ_s depending on original function's derivative $g'(\alpha_s)=\frac{d}{d\alpha_s} g(\alpha_s)$. The division by zero occurs if the difference of two values $g'(\alpha_i)$ numerically equals zero, $\exists i \neq s : g'(\alpha_s)-g'(\alpha_i) = 0$. That is if g resembles a linear function over some interval. We fixed that by removing all basepoints α_s with $g'(\alpha_s)=g'(\alpha_i)$, $i<s$ and inserting these basepoints between basepoints whose $g'(\alpha)$ values differ from each other. This assures that the error never increases or can be reduced: For those intervals of the function which are nearly linear, a linearisation over the whole interval yields a low error; thus, removing basepoints within this interval has little impact. Inserting these basepoints at another non-linear part of the function improves the linearisation accuracy as the PWL function becomes tighter.

3.4.3 Formulation of linearised problem

This section describes the problem reformulation using a PWL function. From existing alternatives [45], we used a *Special Ordered Set* (SOSk) of type $k = 2$ (SOS2) [9]: In a set of continuous variables, at most k of them, adjacent to each other, may take non-zero values. Current linear solvers directly support SOS2.

A PWL function $\tilde{y}=\tilde{g}(x)$ is represented by a set of m continuous decision variables $0 \leq z_s \leq 1$ with a SOS2$(z_0, .. z_s, .. z_{m-1})$ constraint and a convex combination $1=\sum_s z_s$. This way, two adjacent values sum up to $1 = z_s + z_{s+1}$. These values are then used as weights for the basepoints $(\alpha_s,\, \beta_s)$ obtained previously by the

7. This paper's extended version details our improvements of Imamoto's algorithm.

linearisation process (**??**). This way, the weighted sum of all basepoints results in the piece-wise linear problem, $x{=}\sum_s z_s\alpha_s$, $y{=}\sum_s z_s\beta_s$.

Using this representation, we linearise the convex part of the objective function (8) $\Lambda_f \mathsf{T}_\mu(\Lambda_f)$, $\Lambda_f{=}\sum_c x_{cf}$, and substitute it by corresponding weighted basepoint sums, the SOS2 constraint, and a convex combination. First, we focus on one facility location and then add indexes to model all locations. For location f, function T_μ computes the TiS (26). With its linearised version $\tilde{\mathsf{T}}_\mu$ (27) the convex part of the objective function, $\Lambda_f \mathsf{T}_\mu(\Lambda_f)$, becomes $\Lambda_f \sum_s \beta_s z_s$. As Λ_f depends on decision variable x_{cf}, multiplying x_{cf} with z_s turns the replacement term to be quadratic; only function T_μ was linearised, not the whole objective function. However, having a linear and not quadratic objective function would reduce problem complexity and speeds up solving. The quadratic term $\Lambda_f \sum_s \beta_s z_s$ needs to be replaced with an equivalent linear term. This is achieved by "moving" the weight Λ_f into function T_μ, which becomes T^w_μ (28). Using $\tilde{\mathsf{T}}^w_\mu$'s basepoints will transform the quadratic into the linear term $\sum_s \beta_s z_s$; the ordinate basepoints are now already weighted. Since the other parts of the objective function were linear, the whole objective function is now linear.

$$\mathsf{T}_\mu(x) = \frac{1}{\mu{-}x} = y, \text{ with } x{=}\Lambda_f \tag{26}$$

$$\tilde{\mathsf{T}}_\mu : x = \sum_s z_s\alpha_s,\ \tilde{y} = \sum_s z_s\beta_s \tag{27}$$

$$\mathsf{T}^w_\mu(x) = x\mathsf{T}_\mu(x) = \frac{x}{\mu{-}x} = y, \text{ with } x{=}\Lambda_f \tag{28}$$

To model all locations, index f is added for each facility location forming the decision variables z_{fs} and basepoint variables α_{fs}, β_{fs}. The linearised version $\widetilde{\mathsf{QP}}(G, \lambda_c, \mu_f, p, \alpha_s, \beta_s)$ is hence:

$$\min_{x,y,z} \quad \frac{1}{\Lambda}\sum_{cf} x_{cf} l_{cf} + \frac{1}{\Lambda}\sum_{fs} \beta_{fs} z_{fs} \tag{29}$$

$$\text{s.t.} \quad \sum_f x_{cf} = \lambda_c, \qquad \forall c \text{ (demand)} \tag{30}$$

$$\sum_c x_{cf} = \sum_s \alpha_{fs} z_{fs}, \qquad \forall f \text{ (capacity)} \tag{31}$$

$$\sum_s z_{fs} = 1,\ \text{SOS2}(z_{f\cdot,\cdot}), \forall f \qquad \text{(pwl)} \tag{32}$$

$$\sum_c x_{cf} \le y_f, \qquad \forall f \text{ (force flip)} \tag{33}$$

$$\sum_f y_f = p \qquad \text{(limit)} \tag{34}$$

The demand-weighted TiS is represented by term $\sum_s \beta_{fs} z_{fs}$ and the corresponding arrival rate at f is $\sum_s \alpha_{fs} z_{fs}{=}\sum_f \Lambda_f{=}\sum_c x_{cf}$; the new capacity constraint (31). This capacity constraint also implicitly assures that the queuing system is in a steady state through the upper bound of the linearization interval $\alpha_{(m-1)} < \mu$; τ from the old constraint (10) becomes obsolete.

The search space of QP consists of $|F|$ binary and $|F||C|$ real variables. In addition $\widetilde{\mathsf{QP}}$ has $m|F|$ real, restricted SOS variables. If both problems were linear we could guess that solving the second problem $\widetilde{\mathsf{QP}}$ takes longer than QP because the search space is larger. However, linear problems are usually solved faster than non-linear problems. Which problem is solved faster? The answers is not obvious, e.g., $\widetilde{\mathsf{QP}}$ is linear but has a larger search space. Their runtimes are experimentally evaluated in **??**.

The maximal linearisation error (35) of the objective function (29) depends on the errors of the linearised parts, which is the sum of used resources, $y_f{=}1$, and their maximal linearisation errors $\epsilon_{\tilde{T}^w_\mu}$.

$$\frac{1}{\Lambda}\sum_{f,\, y_f=1} \epsilon_{\tilde{T}^w_{\mu_f}} \ \le\ \frac{p}{\Lambda}\max_{\tilde{T}^w}\{\epsilon_{\tilde{T}^w}\} \tag{35}$$

The linearisation accuracy drops if more resources are allowed to open (p). To maintain the same linearisation accuracy while doubling p the linearisation error $\epsilon_{\tilde{T}^w}$ has to be halved. This can be achieved by using more basepoints for the linearisation. Even if **??** indicates that less than twice the basepoints are necessary, increasing the number of basepoints and, hence, increasing the problem's search space increases the runtime.

4 Evaluation

This section has four parts. First, it presents different TiS function linearisations to find a balance of two conflicting goals: Small objective function approximation error and few basepoints (m) for fast computation. Second, solutions of the convex and linear problem (QP vs. $\widetilde{\mathsf{QP}}$) are compared for different real networks. Third, the trade-off between the number of used locations and resulting response time is discussed. Finally, application and network properties are presented for which considering the QD yields better response times than ignoring QD ($\widetilde{\mathsf{QP}}$ vs. P).

4.1 Weighted TiS Linearization

This sections describes how we obtain the concrete basepoints for $\mathsf{T}^w_\mu(\lambda)$ (28) in the evaluation. For this, we show a simplification with one set of basepoints adapted at runtime for different μ values. Afterwards, we discuss the trade-off between a fast solving time and low approximation error.

Function $\mathsf{T}^w_\mu(\lambda)$ depends on μ and needs individual linearizations for different μ; let α^μ_s, β^μ_s be their basepoints (36). Alternatively, function $\mathrm{T}^w(\rho)$ (37) is independent of μ with corresponding basepoints α_s, β_s. Function T^w_μ can be rewritten as $\mathrm{T}^w(\lambda/\mu){=}\mathsf{T}^w_\mu(\lambda)$ (38) and the corresponding basepoints can be rewritten similarly: $\forall s : \alpha^\mu_s{=}\mu\alpha_s$, $\beta^\mu_s{=}\beta_s$.

$$\mathsf{T}^w_\mu(\lambda) = \frac{\lambda}{\mu{-}\lambda} : \quad \sum_s \alpha^\mu_s z_s{=}\lambda,\ \sum_s \beta^\mu_s z_s{=}\mathsf{T}^w_\mu(\lambda) \tag{36}$$

$$\mathrm{T}^w(\rho) = \frac{\rho}{1{-}\rho} : \quad \sum_s \alpha_s z_s{=}\rho,\ \sum_s \beta_s z_s{=}\mathrm{T}^w(\rho) \tag{37}$$

$$\mathrm{T}^w(\frac{\lambda}{\mu}) = \frac{\lambda/\mu}{1{-}\lambda/\mu} : \quad \sum_s \alpha_s z_s{=}\frac{\lambda}{\mu},\ \sum_s \beta_s z_s{=}\mathrm{T}^w(\frac{\lambda}{\mu}) \tag{38}$$

As the ordinate basepoints β_s remain unchanged, the basepoints' approximation error is also not affected. With this handy transformation, we only need to precompute basepoints of T^w instead of basepoint sets of T^w_μ for each different μ in the model, which speeds up the model setup process.

In the remaining section, we investigate the trade-off between a fast solving time and a low-error linearisation.

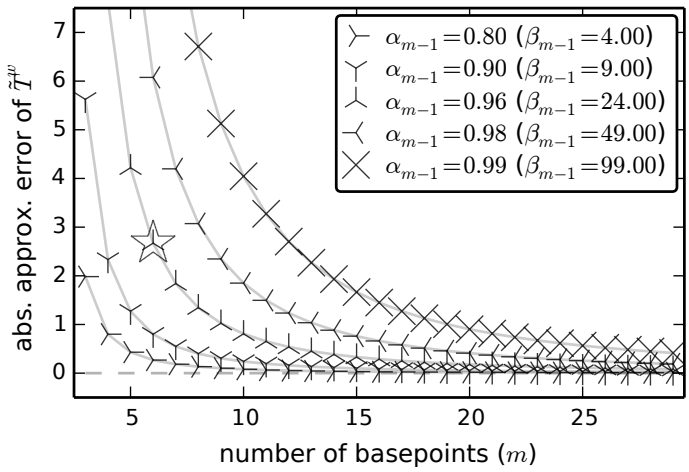

Figure 7: The error of linearising T^w_μ is shown depending on the number of basepoints and different linearization intervals $[0, .., \alpha_{m-1}]$.

For the first, we need to minimize the number of decision variables z_{fs} or, equivalently, the number of basepoints used for the linearisation. For the second, we investigate two control knobs (**??**): many basepoints or small interval end α_{m-1}. **??** shows the error of $\tilde{\mathsf{T}}^w$ depending on the number of basepoints m for different α_{m-1} values. We need a small error (down the vertical axis) with small m (left on the horizontal axis) with large α_{m-1}. The latter also artificially limits the resource capacity and renders solving an input infeasible that could in fact be solved with larger α_{m-1}.

For our evaluation, we set α_{m-1}=0.96 and m=6 with error $\epsilon_{\mathsf{T}^w_\mu}$=2.67 as a good compromise between the number of decision variables, approximation error, and artificial capacity limit.

4.2 Comparison: Convex vs. Linear

We choose the following structurally different topologies from SndLib [44]: `ta2`, `zib54` with many nodes (around 50); `yuan`, `bwin` with few nodes (around 10); `atlanta`, `norway` for dense networks (node:edge ratio 1:2). All topologies are connected. We approximate the latency between nodes by their geographical distance [32]. We assume that data centres are built at well connected nodes/routers, so we selected the 10 nodes with the highest degree[8] to be data centres. We set a relatively low service rate of μ=100 $^{\text{req.}}/_{\text{s}}$ to reflect our computation-intensive example applications [7], [37], [54]. The service rates were the same for all data centres. User requests arrive at all nodes and the arrival rate λ_c for users at site c is randomly generated: Each value is uniformly drawn from an $[0, 1]$ interval and, afterwards, all values are normalized to $\sum_c \lambda_c$=470 $^{\text{req.}}/_{\text{s}}$. This value, together with the service rate μ=100 $^{\text{req.}}/_{\text{s}}$, ensures feasibility for 5 or more facilities. We randomly generated 50 different realizations for all arrival rates. For each of these 50 sets of arrival rates, we considered $p \in [5, 10]$ facilities, resulting in 300 configurations per topology. Each configuration was solved using either QP or $\widetilde{\mathsf{QP}}$.

8. For same degrees the node id is the tie breaker.

?? compares the solution quality (a) and solving time (b). The horizontal axis lists the groups of different topologies and the used number of resources for each group. The vertical axis in (a) shows the 95% confidence intervals of the average response times as the quality of solutions obtained by solving QP (cross) and $\widetilde{\mathsf{QP}}$ (line) for each of the 50 realizations for each group. Similarly, the vertical axis in (b) shows the 95% confidence intervals of the solving time for each group. *The $\widetilde{\mathsf{QP}}$'s solutions have a very similar quality to* QP*'s solutions.*

Looking at (b), obtaining QP's solutions took longer than obtaining the $\widetilde{\mathsf{QP}}$'s solutions. However, those values have to be interpreted with care. Our implementation of QP has to process all possible combinations (F'), whereas $\widetilde{\mathsf{QP}}$ benefits from the MIP solver's branch-and-cut algorithm to reduce the search space. To compare this structurally different problems, we restrict the number of candidate facilities to 10 and the number of possible combinations; the major cause of the higher solving time of QP. Nevertheless, *the absolute solving times of $\widetilde{\mathsf{QP}}$ are very short* for all groups.

In conclusion, *the linearised problem $\widetilde{\mathsf{QP}}$ is an adequate substitution for our original problem* QP*: fast and accurate.*

4.3 Response Time Reduction

?? also shows how response time improves when adding a resource. We could verify two effects decreasing the response times: First, using more locations allows better load balancing, which reduces the queuing delays. These reductions are larger for highly utilized locations than for less utilized ones. Second, more locations allow nearer locations, reducing the round trip times. In conclusion, the average response time of QP$(..., p)$ decreases monotonically in number of resources. Then, service providers earning more money by connecting users with lower response times face a diminishing return. At one breaking point p^* the cost for adding a resource will exceed the additionally earned money. This point depends on the topology, service times, and service monetization. By using QP with different p values, the service provider can determine p^* in advance to avoid profit loss.

For a closer look, **??** shows not only the average response time but also the time in system and round trip times grouped along the horizontal axis the same way like in **??**. We can trace the two effects of response time reductions: First, load balancing across more resources reduces the queuing delay and with it the time in queuing system. Second, the round trip time is reduced as nearer resources are used.

4.4 Considering Queueing Delays

This paper presents a refinement of the assignment problem (P in (4)) by additionally considering queuing delays. This refinement increases problem's complexity, accuracy, and solving time. Only a significantly lower response time would make these drawbacks worthwhile; for instances, small queuing delays compared to large round trip times will render the refinement unnoticeable. This section investigates multiple scenario factors influencing queuing delays and judges the refinement gain by comparing solutions' response times obtained by either ignoring or considering queuing delays.

For this, we perform a second experiment with a slightly different configuration as in **??**. The first experiment showed

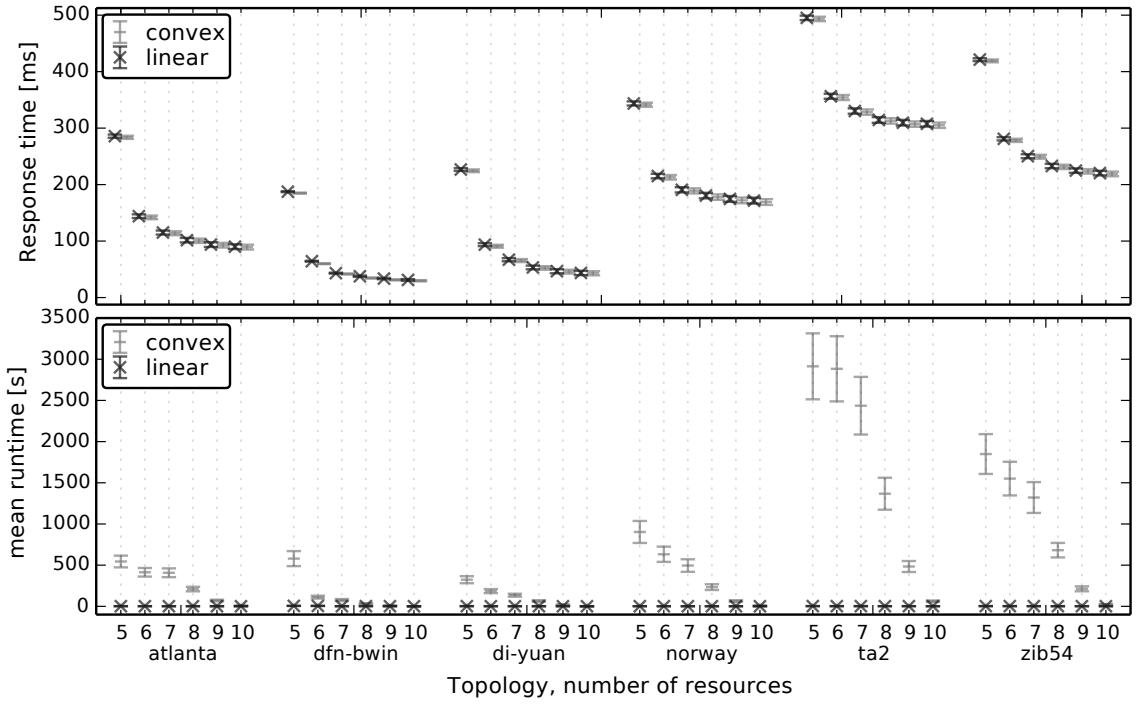

Figure 8: Solution quality (a) and solving time (b) for convex QP and its linearizations $\widetilde{\mathsf{QP}}$. Estimated mean confidence intervals at 95% confidence level.

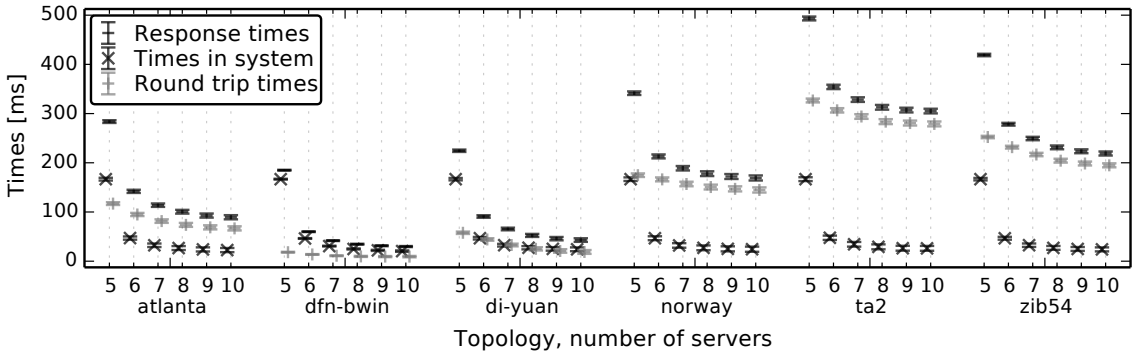

Figure 9: Round trip times, times in system, and round trip times for different topologies and numbers of used servers. Estimated mean confidence intervals at 95% confidence level.

that $\widetilde{\mathsf{QP}}$ is a very accurate and fast substitution for QP. Focussing on $\widetilde{\mathsf{QP}}$ allows us to evaluate more scenario variations in a reasonable time than would be possible with QP.

4.4.1 Configuration

We vary configurations by four factors[9] $\hat{\mu}$, $\hat{D}$, $\hat{\rho}_s$, $\hat{G}$. The service rate $\hat{\mu}$ reflects how computation-intensive the application is. A small service rate means a high processing time. For the same arrival rate, queuing delays are higher for smaller service rates. For clarity, we assume homogeneous[10] resources, same service times at each location, $\hat{\mu}{=}\mu_f$, $\forall f$. Different levels of $\hat{\mu}$ represents different application types ranging from fast web servers with a short processing time up to computation-intensive applications, $\hat{\mu}{=}10{,}000; 1000; 100; 10; 1\,{}^{\text{req.}}/_{\text{s}}$.

Unlike the service rates, the arrival rates λ_c are not homogeneous but randomly distributed. This enables us to investigate different patterns for spatially distributed load, e.g., fluctuations or local hot spots. Let $\hat{\lambda}$ be the targeted mean arrival rate. For this, we choose three different random distributions $\hat{D}$ for our second factor. First, a similar load across all nodes with small fluctuations is represented by a "narrow" normal distribution: $\mathrm{N}(\text{mean}{=}\hat{\lambda}, \text{std.dev.}{=}{}^{\hat{\lambda}}/_{20}){=}\mathrm{N}_1$. Second, a largely fluctuating load around an average load per node is represented by a "wide" normal distribution: $\mathrm{N}(\hat{\lambda}, \hat{\lambda}){=}\mathrm{N}_2$. Third, heavy variations causing local hot spots are represented by an exponential distribution: $\mathrm{Exp}(\hat{\lambda})$. For each node the arrival rate λ_c is drawn from $\hat{D}$, where negative values are capped to

9. The $\hat{}$ indicates the factors under investigation.

10. We discard an additional minimization potential of heterogeneous resources to simplify the comparison: When assigning demand, a farer but faster resource enables trading off a larger RTT for a smaller QD+PT. By doing so, the response time is further reduced.

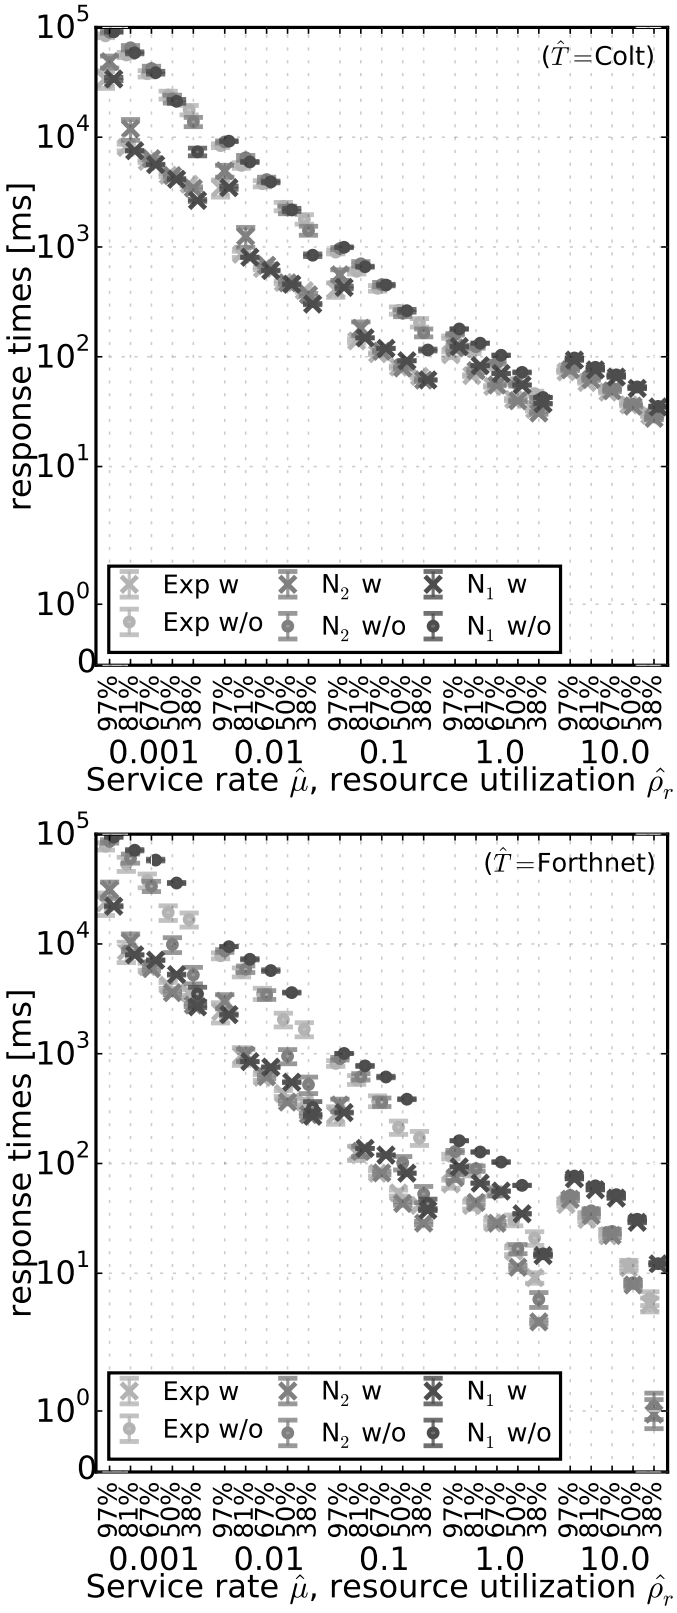

Figure 10: Average response times as a function of factor combination ($\hat{\mu}$, $\hat{\rho}_s$) for different demand distributions ($\hat{D}$) with 95% confidence intervals. Left (and right) comparison for topology $\hat{T}$=`Colt` (and `Forthnet`, respectively).

zero, $\lambda_c = \max\{0, X\},\ X \sim \hat{D} \in \{\mathrm{N}_1, \mathrm{N}_2, \mathrm{Exp}\}$. We investigate 50 different such realizations for each topology.

The third factor $\hat{\rho}_r$ reflects the average resource utilization. Highly utilized resources have high QDs, and solutions are very similar whether or not the queuing delay is considered. To enforce the $\hat{\rho}_r$ levels, we limit the number of utilised resources p defined later.

The last factor, with different topologies $\hat{G}$ we represents structural differences like the diameter or ratio between the round trip time and queuing delays, for instance, sparse graphs have a larger diameter and higher round trip times than dense graphs. We selected 14 out of 524 topologies from different sources: sndlib[11] [44], topology zoo**??** [35], kingtrace[12] [26]. The selection considers three categories for the number of nodes, edge, and diameter to eliminate roughly similar topologies.

In each topology, the 100 best connected nodes[13] were marked as candidate resource locations F, $|F| = \min\{|N|, 100\}$. Let $p_{\min} \leq |F|$ be the number of resources at least necessary to handle all demand. Having more demand using more resources means less freedom for location choices: for instance, if $p_{\min} = |F|$ all resources are fully utilized and only this decision is possible. We set $p_{\min} = \lfloor 0.3|F| \rfloor$ to allow enough freedom. Then, $\hat{D}$'s target arrival rate $\hat{\lambda}$ is defined accordingly, $\hat{\lambda} = \hat{\mu} \frac{p_{\min}}{|N|}$. Actually using $p = p_{\min}$ resources results in a very high resource utilization p_r; allowing more resources reduces the utilization. For our third factor, the resource utilization $\hat{\rho}_r = 0.97; 0.81; 0.67; 0.5; 0.375$, we set corresponding p values achieving (roughly) the targeted server utilization, $\hat{p} = \lfloor {}^{0.3}/_{\hat{\rho}_r} |F| \rfloor = \lfloor a|F| \rfloor$, $a = 0.31; 0.37; 0.45; 0.6; 0.8$.

4.4.2 Results

Summarising, each of the 1050 combinations of the four factors ($\hat{\mu}$, $\hat{D}$, $\hat{\rho}_s$, $\hat{T}$) was randomised by 50 demand realisations, resulting in $52{,}500$ different configurations for which problem P (without considering queuing delays) and problem $\widetilde{\mathsf{QP}}$ (with considering queuing delays) are solved. The quality metric is the average response time computed by QP's exact objective function (8). The difference between response time obtained with $\widetilde{\mathsf{QP}}$ and with P measures the response time improvement when considering queuing delays. A larger difference means $\widetilde{\mathsf{QP}}$'s assignments are superior to P's assignments. As P is a simplification of QP, P's response times (RTs) cannot be smaller than QP's RTs. Only using the linearisation $\widehat{\mathsf{QP}}$ could potentially result in a higher RT, but never occur[14] in our evaluation.

?? shows response times as a function of two factors, service rate $\hat{\mu}$ and resource utilisation $\hat{\rho}_s$ for two selected topologies $\hat{T}$=`Colt`, `Forthnet`. Each single data point corresponds to the average response times with 95% confidence intervals of 50 realisations for one factor combination $(\hat{\mu}, \hat{\rho}_s, \hat{T})$. At a first glance, the response times of the two problems can be compared column-wise: The circle data points represent problem P's solutions. Most of them are above the cross data points representing the $\widetilde{\mathsf{QP}}$'s solutions, no circle appears above a cross: The difference between a column-wise cross circle pair visualizes the response time improvements when considering queuing delays; the y-axis is in log-scale.

At a second look, **??** shows another pattern along increasing service rates ($\hat{\mu}$) from $1\,{}^{\mathrm{req.}}/_{\mathrm{ms}}$ to $10{,}000\,{}^{\mathrm{req.}}/_{\mathrm{ms}}$ (left to right): Significant response time improvements for low service rates and nearly no response time improvements for high service rates. *When queuing delay (and processing times) become a significant part of the response time, it is necessary to consider queuing delays when deciding the assignments* as otherwise unnecessary higher response times would be the result.

At a third look, **??** shows larger response time differences for higher utilized ($\hat{\rho}_s$) configurations; in these cases, the queuing delay becomes a dominant part of the response time. *Considering queuing delays is necessary for utilised topologies ($\rho_s > 0.5$).*

At a fourth look, **??** shows how the demand distributions represented by different grey levels influence the response time: While the `Colt` topology show few changes in response times for different demand distributions, the `Forthnet` topology data points are more spread for different demand distributions. In tendency, the more challenging demand distributions (Exp and N_2) show larger response time improvements than the other distribution (N_1). Considering queuing delays avoids assigning local demand hotspots to nearby but highly utilized resources; instead more and lesser utilized resources are used effectively reducing the average response time.

In summary, *the evaluation verifies our claim, that considering queuing delays is necessary as response times can significantly reduced.* This applies only when queuing delays and round trip times are of the same magnitude; e.g. configuration with service rate $\hat{\mu} = 10{,}000\,{}^{\mathrm{req.}}/_{\mathrm{ms}}$ (short queuing delays) have only marginal response time improvements when considering queuing delays.

5 Conclusion

We extend previous work by optimally solving the assignment problem QP, a Facility Location Problem with integrated queuing systems. We proposed problem $\widetilde{\mathsf{QP}}$ as a linearisation of QP with accurate solutions obtained fast. In our scenario of adapting the resource allocation at geographically distributed sites, such a swift reaction is important. It allows to swiftly react immediately to an ever changing environment including demand fluctuations or network congestion.

We showed that adding more and more resources will at one point reduce the user expected response time only marginally. With our work, the application provider can determine this point in advance and can allocate resources accordingly.

11. Round trip times were approximated by geographical distances [32].

12. A sparse matrix specifies point to point latencies. Some were only available in one direction. We assume the same latency for the opposite direction; otherwise those nodes had to be discarded.

13. Highest degree first; node ID as the tie breaker.

14. Some pairs of $\widetilde{\mathsf{QP}}$ and P solutions are treated as equivalent good when their response time differences were below 10^{-6} – the solver threshold at which the solver stopped improving the solution in our evaluation.

We performed a large-scale experiment and traced down network and application properties where integrating the queuing system into the FLP improves solutions. This could guide other researchers or application providers whether the complex problem QP is necessary to apply or the simpler problem P is sufficient enough for their scenario.

The simple M/M/1-queue model can be replaced with more sophisticated queuing models as long as (i) the inter-arrival times are described by a Poisson process and (ii) the queuing delay function is convex. For models with a different inter-arrival time process, splitting and joining becomes much more complicated. The linearised problem even supports non-convex queuing delay functions, but for such models the presented algorithm for obtaining basepoints is no more applicable.

Finally, the shown modelling techniques can be used beyond our use case. The assignment problem QP is part of the broader family of *FLPs with convex cost functions*. We think for most, maybe all of them, a similar good and fast problem linearisation can be formulated by reusing our problem linearization formulation and by determining the PWL basepoints with the presented algorithm.

REFERENCES

[1] Amazon Web Service.

[2] R. Aboolian, O. Berman, and Z. Drezner. The multiple server center location problem. *Annals of Operations Research*, 167(1):337–352, mar 2009.

[3] S. Agarwal, J. Dunagan, and N. Jain. Volley: Automated data placement for geo-distributed cloud services. In *Proceedings of the 7th conference on Networked systems design and implementation (NSDI '10)*. USENIX, 2010.

[4] M. Alicherry and T. Lakshman. Network aware resource allocation in distributed clouds. In *2012 Proceedings IEEE INFOCOM*, pages 963–971. IEEE, mar 2012.

[5] M. S. Andersen and L. Vandenberghe. CVXOPT - a free software package for convex optimisation, 2013.

[6] M. Bagaa, T. Taleb, and A. Ksentini. Service-aware network function placement for efficient traffic handling in carrier cloud. In *Proceedings of the Wireless Communications and Networking Conference (WCNC)*, pages 2402–2407. IEEE, apr 2014.

[7] S. K. Barker and P. Shenoy. Empirical evaluation of latency-sensitive application performance in the cloud. In *Proceedings of the 1st annual conference on Multimedia systems (MMSys)*, New York, 2010. ACM Press.

[8] M. Bauer, S. Braun, and P. P. Domschitz. Media Processing in the Future Internet. In *Proceedings of the 11th Würzburg Workshop on IP: Visions of Future Generation Networks*, pages 113–115, Würzburg, 2011.

[9] E. Beale, M. Lansdowne, and J. A. Tomlin. Special facilities in a general mathematical programming system for non-convex problems using ordered sets of variables. *Operational Research*, 69(447-454):99, 1970.

[10] O. Berman and Z. Drezner. The multiple server location problem. *Journal of the Operational Research Society*, 58(1):91–99, 2007.

[11] G. Bolch, S. Greiner, H. De Meer, and K. S. Trivedi. *Queueing networks and Markov chains: Modeling and Performance Evaluation with Computer Science Applications*. Wiley-Interscience, 2005.

[12] S. Boyd and L. Vandenberghe. *Convex optimization*, volume 25. Cambridge University Press, jun 2004.

[13] D. Cai and S. Natarajan. The Evolution of the Carrier Cloud Networking. In *Seventh International Symposium on Service-Oriented System Engineering*, pages 286–291. IEEE, mar 2013.

[14] V. Cardellini, M. Colajanni, and P. Yu. Geographic load balancing for scalable distributed Web systems. In *Proceedings of the 8th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, pages 20–27. IEEE Comput. Soc, 2000.

[15] K. Church, A. Greenberg, and J. Hamilton. On Delivering Embarrassingly Distributed Cloud Services. In *HotNets*, 2008.

[16] M. Claypool and K. Claypool. Latency and player actions in online games. *Communications of the ACM - Entertainment networking*, 49(11):40–45, 2006.

[17] A. Csaszar, W. John, M. Kind, C. Meirosu, G. Pongracz, D. Staessens, A. Takacs, and F. J. Westphal. Unifying cloud and carrier network: EU FP7 Project UNIFY. In *6th International Conference on Utility and Cloud Computing*, pages 452–457. IEEE/ACM, 2013.

[18] T. Cucinotta, K. Oberle, M. Stein, P. Domschitz, and S. Mullender. Run-time Support for Real-Time Multimedia in the Cloud. In *2nd International Workshop on Real-Time and Distributed Computing in Emerging Applications (REACTION 2013)*, 2013.

[19] T. Drezner and Z. Drezner. The gravity multiple server location problem. *Computers & Operations Research*, 38(3):694–701, mar 2011.

[20] Z. Drezner and H. W. Hamacher. *Facility location: applications and theory*. Springer, 2004.

[21] P. T. Endo, A. de Almeida Palhares, N. Pereira, G. Goncalves, D. Sadok, J. Kelner, B. Melander, and J.-E. Mangs. Resource allocation for distributed cloud: concepts and research challenges. *IEEE Network*, 25(4):42–46, 2011.

[22] A. Fischer, J. F. Botero, M. T. Beck, H. de Meer, and X. Hesselbach. Virtual Network Embedding: A Survey. *Communications Surveys & Tutorials, IEEE*, (99):1–19, 2013.

[23] M. J. Freeman, K. Lakshminarayanan, and David Mazieres. OASIS: Anycast for Any Service. In *Networked Systems Design and Implementation (NSDI)*, 2006.

[24] B. Geissler, A. Martin, A. Morsi, and L. Schewe. Using Piecewise Linear Functions for Solving MINLPs. In J. Lee and S. Leyffer, editors, *Mixed Integer Nonlinear Programming*, volume 154 of *The IMA Volumes in Mathematics and its Applications*, pages 287–314. Springer New York, 2013.

[25] H. Goudarzi and M. Pedram. Geographical Load Balancing for Online Service Applications in Distributed Datacenters. In *6th International Conference on Cloud Computing*, pages 351–358. IEEE, 2013.

[26] K. P. Gummadi, S. Saroiu, and S. D. Gribble. King: Estimating Latency between Arbitrary Internet End Hosts. In *Proceedings of the 2nd Workshop on Internet measurment (IMW '02)*, pages 5–18. ACM SigComm, 2002.

[27] O. Hakimi and S. L. Kariv. An Algorithmic Approach to Network Location Problems. II: The p-Medians. *SIAM Journal on Applied Mathematics*, 37(3):539–560, feb 1979.

[28] F. Hao, T. V. Lakshman, S. Mukherjee, and H. Song. Enhancing dynamic cloud-based services using network virtualization. *Proceedings of the 1st ACM workshop on Virtualized infrastructure systems and architectures (VISA)*, 40(1):37, 2009.

[29] A. Imamoto and B. Tang. Optimal Piecewise Linear Approximation of Convex Functions. In *World Congress on Engineering*, pages 22—-25, 2008.

[30] A. Ishii and T. Suzumura. Elastic Stream Computing with Clouds. In *Proceedings of the 4th International Conference on Cloud Computing*, pages 195–202. IEEE, jul 2011.

[31] O. Kariv and S. L. Hakimi. An algorithmic approach to network location problems. The p-medians. *SIAM Journal on Applied Mathematics*, 37(3):539–560, 1979.

[32] S. Kaune, K. Pussep, C. Leng, A. Kovacevic, G. Tyson, and R. Steinmetz. Modelling the internet delay space based on geographical locations. In *Proceedings of the 17th Euromicro International Conference on Parallel Distributed and Network-based Processing*, pages 301–310. IEEE, 2009.

[33] M. Keller and H. Karl. Response Time-Optimized Distributed Cloud Resource Allocation. In *Workshop on Distributed Cloud Computing (DCC)*. ACM, 2014.

[34] M. Keller, M. Peuster, C. Robbert, and H. Karl. A Topology-aware Adaptive Deployment Framework for Elastic Applications. In *17th International Conference on Intelligence in Next Generation Networks (ICIN)*, pages 61–69. IEEE, oct 2013.

[35] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan. The Internet Topology Zoo. *Journal on Selected Areas in Communications*, 29(9):1765–1775, oct 2011.

[36] M. Kwon. A Tutorial on Network Latency and its Measurements. *IGI Global*, 2014.

[37] Y.-T. Lee, K.-T. Chen, H.-I. Su, and C.-L. Lei. Are all games equally cloud-gaming-friendly? an electromyographic approach. In *Proceedings of the 11th Annual Workshop on Network and Systems Support for Games (NetGames)*, pages 4–9. ACM/IEEE, 2012.

[38] M. Lin, Z. Liu, A. Wierman, and L. L. H. Andrew. Online algorithms for geographical load balancing. In *International Green Computing Conference (IGCC)*, pages 1–10, 2012.

[39] Z. Liu, M. Lin, A. Wierman, S. Low, and L. Andrew. Greening Geographical Load Balancing. *Transactions on Networking (TON)*, 23(2):657–671, 2015.

[40] V. Marianov and D. Serra. Probabilistic maximal covering location-allocation models with constrained waiting time or queue length for congested systems. *Journal of Regional Science*, 38(3):401–424, 1996.

[41] V. Marianov and D. Serra. Location – Allocation of Multiple-Server Service Centers. *Annals of Operations Research*, 111(1–4):35—-50, 2002.

[42] M. Moghadas and T. Kakhki. Maximal covering location-allocation problem with M/M/k queuing system and side constraints. *Iranian Journal of Operations Research*, 2(2):1–16, 2011.

[43] K. Oberle, D. Cherubini, and T. Cucinotta. End-to-end service quality for cloud applications. In *Economics of Grids, Clouds, Systems, and Services*, volume 8193 LNCS, pages 228–243. Springer, 2013.

[44] S. Orlowski, M. Pioro, A. Tomaszewski, and R. Wessaly. SNDlib Survivable Network Design Library. Technical Report June, Konrad-Zuse-Zentrum für Informationstechnik Berlin, Spa, Belgium, 2007.

[45] M. Padberg. Approximating separable nonlinear functions via mixed zero-one programs. *Operations Research Letters*, 27(1):1–5, aug 2000.

[46] S. Pandey, A. Barker, K. K. Gupta, and R. Buyya. Minimizing Execution Costs when Using Globally Distributed Cloud Services. In *Proceedings of the 24th International Conference on Advanced Information Networking and Applications*, pages 222–229. IEEE, 2010.

[47] S. H. R. Pasandideh and S. T. A. Niaki. Genetic application in a facility location problem with random demand within queuing framework. *Journal of Intelligent Manufacturing*, 23(3):651–659, may 2010.

[48] M. Scharf, T. Voith, W. Roome, B. Gaglianello, M. Steiner, V. Hilt, and V. K. Gurbani. Monitoring and abstraction for networked clouds. In *Proceedings of the 16th International Conference on Intelligence in Next Generation Networks (ICIN)*, pages 80–85. IEEE, oct 2012.

[49] B. Shah and B. H. Trivedi. Improving Performance of Mobile Agent Based Intrusion Detection System. *Fifth International Conference on Advanced Computing & Communication Technologies*, pages 425–430, 2015.

[50] A. L. Stolyar and B. Labs. Shadow-Routing Based Dynamic Algorithms for Virtual Machine Placement in a Network Cloud. In *Proceedings of the 32nd International Conference on Computer Communications (INFOCOM '13)*, pages 644–652. IEEE, 2013.

[51] T. Taleb. Toward carrier cloud: Potential, challenges, and solutions. *Wireless Communications, IEEE*, 21(3):80—-91, jun 2014.

[52] T. Taleb and A. Ksentini. Follow Me cloud: Interworking federated clouds and distributed mobile networks. *IEEE Network*, 27(5):12–19, 2013.

[53] N. Vidyarthi, S. Elhedhli, and E. Jewkes. Response time reduction in make-to-order and assemble-to-order supply chain design. *IIE Transactions*, 41(5):448–466, mar 2009.

[54] Z. Wan. Cloud Computing infrastructure for latency sensitive applications. In *Proceedings of the 12th International Conference on Communication Technology*, pages 1399—-1402. IEEE, nov 2010.

[55] Q. Wang, R. Batta, and C. M. Rump. Algorithms for a Facility Location Problem with Stochastic Customer Demand and Immobile Servers. *Annals of Operations Research*, 111(1-4):17–34, mar 2002.

[56] P. Wendell, J. W. Jiang, M. J. Freedman, and J. Rexford. DONAR : Decentralized Server Selection for Cloud Services. In *SigComm*, volume 40, pages 231–242, 2010.

[57] B. Wong. ClosestNode. com: an open access, scalable, shared geocast service for distributed systems. *ACM SIGOPS Operating Systems Review*, pages 62–64, 2006.

[58] B. Wong, A. Slivkins, and E. Sirer. Meridian: A lightweight network location service without virtual coordinates. In *ACM SIGCOMM Computer Communication Review*, volume 35, pages 85–96. ACM, 2005.

[59] Q. Zhang, Q. Zhu, M. F. Zhani, and R. Boutaba. Dynamic Service Placement in Geographically Distributed Clouds. In *Proceedings of the 2nd International Conference on Distributed Computing Systems*, pages 526–535. IEEE, jun 2012.

Matthias Keller received his diploma degree in computer science from the University of Paderborn, Germany. He is currently working as a research associate in Computer Networks group in University of Paderborn. He focuses on adaptive resource allocation across wide area networks, designed a framework, and created a prototype testbed. Previously, he worked at the Paderborn Centre for Parallel Computing as a research associate and at an international software development company as a software engineer.

Holger Karl received his PhD in 1999 from Humboldt University Berlin; afterwards he joined Technical University Berlin. Since 2004, he is Professor for Computer Networks at University Paderborn. He is also responsible for the Paderborn Centre for Parallel Computing and has been involved in various European and national research projects. His main research interests are wireless communication and architectures for the Future Internet.